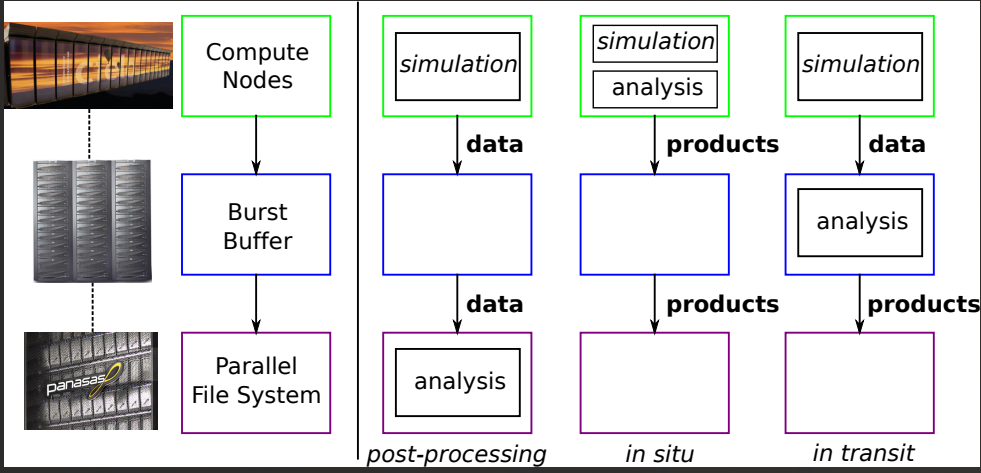


Future Technologies in Computing at Los Alamos



Above: New data-analysis workflows including in situ and in transit methods are well-suited to burst-buffer file systems

Extreme-Scale Burst-Buffer Storage Systems

Large high-performance computing (HPC) systems push parallel file systems to extremes in aggregate I/O bandwidth, and numbers of clients. Current HPC applications generate I/O in bursts driven by algorithmic and checkpoint/restart needs, so an attractive approach for meeting future I/O needs is to adopt a tiered storage system design that integrates layers of solid-state burst buffers to absorb application requests and deliver a lower-bandwidth, sustained load to the base file system. Los Alamos is leading efforts to develop burst buffer designs, characterize their performance and benefits, and to develop software interfaces to isolate application developers from hardware implementation details. In many attractive designs, processors are integrated into the burst buffer; offering tantalizing opportunities to move computation into the I/O stream for in transit analysis, indexing, and data reduction. The challenges of software and scheduling for such systems are being actively explored in a number of Los Alamos' projects.

Contacts: Gary Grider, ggrider@lanl.gov; Jonathan Woodring, woodring@lanl.gov

Sponsor: DOE NNSA/ASC and LANL LDRD

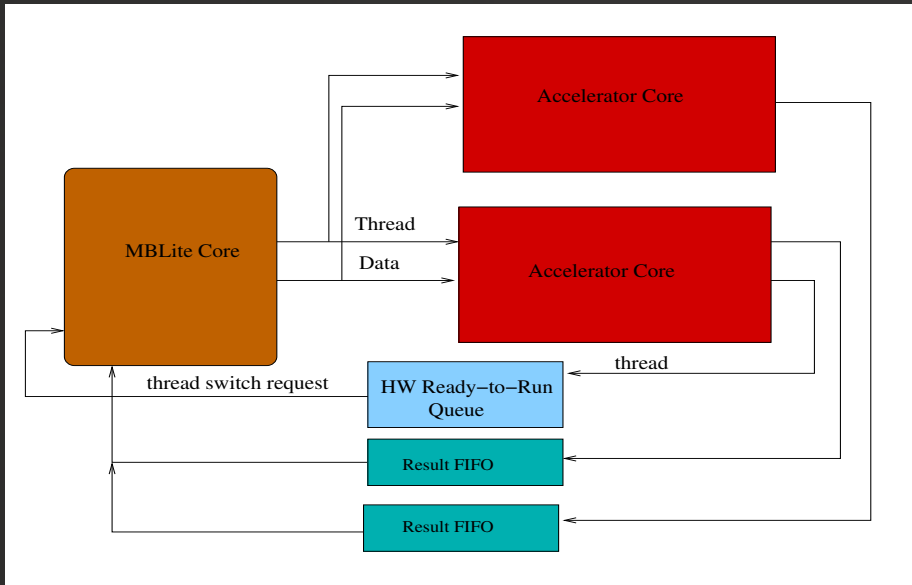
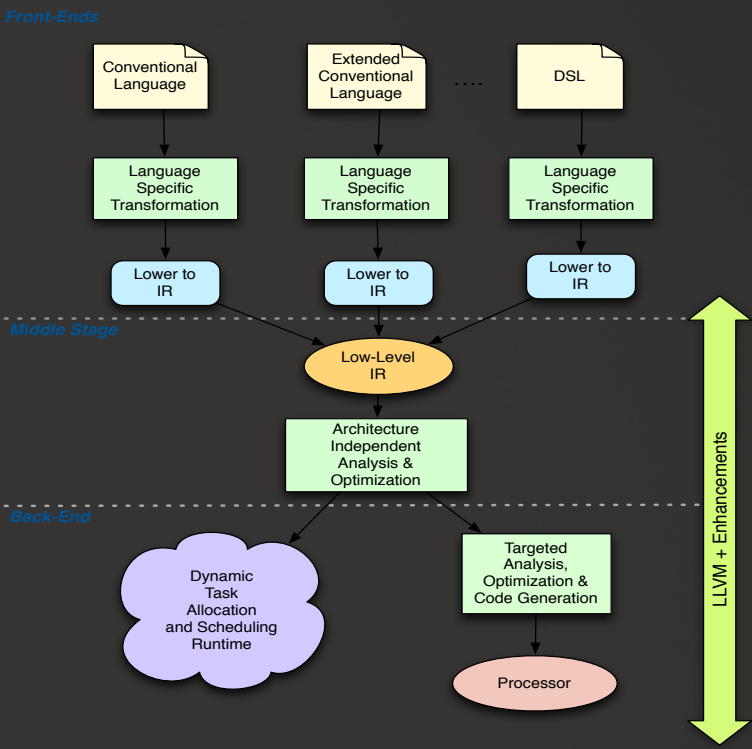
Open-Source Compilers and Runtimes for Emerging Computer Architectures

New architectures require software that explicitly targets shared and distribute-memory parallelism and deep memory hierarchies.

We are addressing these challenges by developing techniques for supporting multiple Domain Specific Languages and language extensions in a single compiler framework, together with a runtime system that performs three critical functions: scheduling of tasks on processing resources, scheduling data access and motion as needed by the tasks, and just-in-time compilation for dynamic optimization of task code to match scheduling constraints.

Contact: Patrick McCormick, pat@lanl.gov

Sponsor: DOE NNSA/ASC and ASCR



Above: A schematic representation of an FPGA design for high-performance computing

Custom Hardware for Scientists: OpenCL to FPGA

Field-Programmable Gate Arrays (FPGAs) offer excellent performance for custom hardware designs with high-parallelism and I/O density and superior energy efficiency. They have been used widely at Los Alamos in application-specific contexts such as space-based computing. The high engineering effort to achieve performance, and limited flexibility for minor changes, however, have meant they have had little impact in general-purpose high-performance scientific computing.

A promising approach for overcoming these drawbacks is the use of direct OpenCL to FPGA hardware compilation. OpenCL is an open standard for cross-platform parallel programming that allows complex kernels to be prototyped on emerging architectures (such as GPUs), and also to be transitioned from traditional architectures to FPGAs. Current applications include N-body force calculations, multigrid solvers, and hydrodynamics with adaptive mesh refinement.

Contact: Zachary Baker, zbaker@lanl.gov

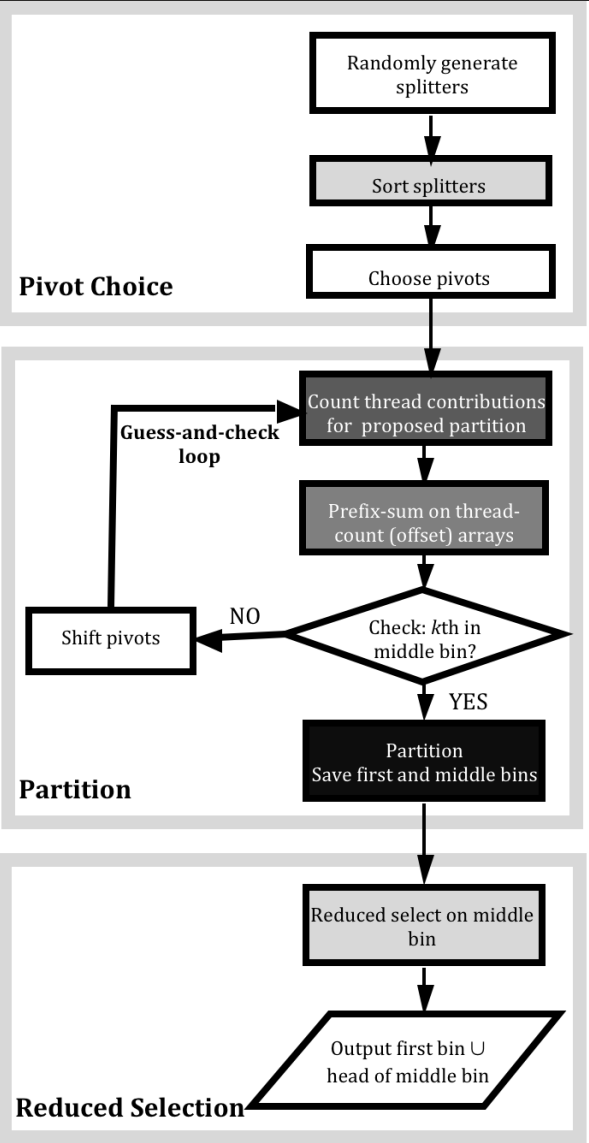
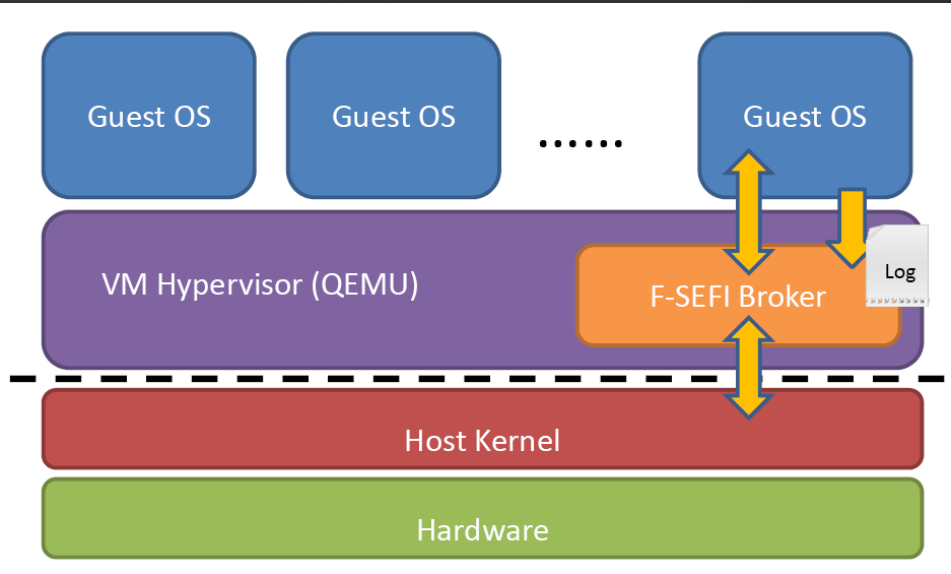
Sponsor: DOE NNSA/ASC

Studying Application Resilience Through Software Fault Injection

Shrinking processor technology, growing supercomputer size, and power-conserving near-threshold-voltage techniques all pose major challenges for high-performance computer system reliability. We are exploring the implications through software fault injection based upon a robust processor emulator virtual machine (QEMU). The virtual machines are run in batch mode on supercomputers with external control of fault injection in specific locations of an application. We can target specific applications, specific functions of that application, and inject faults in extremely configurable ways. Then, we observe how the application responds to this fault and explore application resilience techniques to address vulnerabilities.

Contact: Nathan DeBardeleben, ndebard@lanl.gov

Sponsor: DOE NNSA/ASC and DOD



Probabilistic Computing for New Computer Architectures

Probabilistic computing encompasses probabilistic hardware or randomized software methods, or both. Probabilistic computation either does not give a deterministic result, or else calculates along non-deterministic paths to reach a deterministic result.

Probabilistic computing is both a challenge and an opportunity: a challenge, since reductions in feature size are introducing increasing non-determinism, and an opportunity to explore a richer space of algorithms, and since we may be able to reduce power consumption by certain probabilistic hardware methods.

We are currently researching probabilistic computational methods on emerging architectures, and have achieved speedups over traditional methods on basic algorithms.

Contact: Laura Monroe, lmunroe@lanl.gov

Sponsor: LANL LDRD and UC Fees Program

Left: An example of a probabilistic selection algorithm being examined in our advanced architecture research